

An Evaluation of Three-Dimensional Sensors for the  
Extravehicular Activity Helper/Retriever

Final Report

NASA/ASEE Summer Faculty Fellowship Program - 1993

Johnson Space Center

Prepared by:	Michael Magee, Ph.D.
Academic Rank:	Professor
University & Department:	University of Wyoming Computer Science Department Laramie, Wyoming 82071-3682
NASA/JSC	
Directorate:	Engineering
Division:	Automation and Robotics
Branch:	Intelligent Systems
JSC Colleague:	Thomas W. Pendleton
Date Submitted:	August 6, 1993
Contract Number:	NGT-44-001-800

## ABSTRACT

The Extravehicular Activity Retriever/Helper (EVAHR) is a robotic device currently under development at the NASA Johnson Space Center that is designed to fetch objects or to assist in retrieving an astronaut who may have become inadvertently de-tethered. The EVAHR will be required to exhibit a high degree of intelligent autonomous operation and will base much of its reasoning upon information obtained from one or more three-dimensional sensors that it will carry and control. At the highest level of visual cognition and reasoning, the EVAHR will be required to detect objects, recognize them, and estimate their spatial orientation and location. The recognition phase and estimation of spatial pose will depend on the ability of the vision system to reliably extract geometric features of the objects such as whether the surface topologies observed are planar or curved and the spatial relationships between the component surfaces. In order to achieve these tasks, accurate sensing of the operational environment and objects in the environment will therefore be critical. This report documents the qualitative and quantitative results of empirical studies of three sensors that are capable of providing three-dimensional information to the EVAHR, but using completely different hardware approaches. The first of these devices is a phase shift laser with an effective operating range (ambiguity interval) of approximately 15 meters. The second sensor is a laser triangulation system designed to operate at much closer range and to provide higher resolution images. The third sensor is a dual camera stereo imaging system from which range images can also be obtained. The remainder of the report characterizes the strengths and weaknesses of each of these systems relative to quality of data extracted and how different object characteristics affect sensor operation.

## PERCEPTRON LASER RANGE SCANNER

### Operational Characteristics

The Perceptron Laser Range Scanner measures distances based on the phase shift of a modulated signal carried on an infrared laser beam. The range values returned by the scanner are represented by 12 bit integers that span a single ambiguity interval of approximately 15.2 meters. This means that a difference of one range unit (out of 4096) represents a distance change of about 4 mm. The scanner is able to produce a dense range image by employing a rotating mirror whose rotation axis can be tilted. The scanner simultaneously provides two separate range and reflectance (intensity) images that are fully registered.

### Observed Characteristics of the Perceptron Laser Scanner

The quality of the range data provided by the scanner is affected by several factors which generally relate to the composition of the surface material, its reflectivity characteristics, its geometry, and the orientation of surface normals relative to the scanner itself. Without question, the most influential among these factors was the reflectivity of the surface material. For extreme cases in which a scanned region was made of a highly specularly reflective material, no reliable range estimate would be expected since the laser beam would be reflected away from the the sensor. This expectation was verified in the case of the Perceptron.

For less extreme cases involving diffuse reflective surfaces, however, the quality of the data was highly dependent on the albedo of the surface. These dependencies can best be illustrated by examining the quality of the range images acquired by scanning black and white planar surfaces (sheets of paper) that were oriented perpendicular to the optical axis of the scanner. As a measure of data stability, the local standard deviation ( $\sigma$ ) for range values was computed within a row. This local standard deviation was based on the center range value and the nearest 8 neighbors within the row. It was observed that the local  $\sigma$  varied by as much as 3 range units. For such cases, in excess of 99% of the range samples could be expected to fall within 3  $\sigma$  ( $\pm 9$  range units) of the mean value. For the test case under discussion, this translates into a local variation of approximately  $\pm 33$  mm over a distance of 8 mm. For the black surface, the quality of the data was significantly worse. Local standard deviations as high as 9 range

values were observed meaning that a 3 sigma test would include range values as far as  $\pm 100$  mm over this limited region of a scan line. The local standard deviations for reflectances varied up to 30 units for the white surface and up to 8 units for the black surface.

The implications of these observed local variations are very important when designing algorithms that attempt to segment the image into component regions such as planes and curved surfaces. For example, the magnitude of the local variations in range values makes it extremely difficult to segment planar surfaces based on a local geometric constraint such as surface normal consistency. Furthermore, even on white objects, it is difficult to recognize the curvature of objects smaller than 100 mm since the magnitude of local range variation is large relative to surface size. If the data is smoothed by a classical filtering mechanism, finer details that are necessary to recognize an object and/or estimate its pose may be lost. Hence, algorithms that depend on local geometry are less likely to succeed than those that take a more global approach to object analysis. The results of both local and global algorithms that were developed are presented in the next section.

### Finding Planes, Recognizing Objects and Estimating their Spatial Poses

The local instability of range values observed for the laser scanner makes scene segmentation using locally computed surface normals difficult unless the range values are smoothed using a reasonably large filter. Applying such a filter, of course, results in a loss of scene detail but does make it possible to find planes that are large relative to the size of the filter.

A better approach to finding surfaces was to grow them based on local range and reflectance difference constraints. It was determined that after applying a 7X7 mean filter, planes that were not highly oblique to the sensor axis could be successfully grown by adding to regions neighboring image elements whose smoothed reflectance and range values did not differ by more than 40 and 1.5, respectively. This provided the basis by which planar regions could be segmented and used as the basis for recognizing one of the Orbital Replacement Units (ORU) and estimating its pose using the normals of three orthogonal surfaces and a point on the handle. This feature matching method employed was based on the computed surface areas and is therefore subject to be sensitive to occlusion. Once the surfaces have been grown, however, other features could be computed such

as vertices from the intersections of planes. A large number of such features (more than four) would provide the basis for feature matching and pose estimation since there would be an overdetermined set of model/image point pairs for computing pose.

The fundamental problem with the above approach is that region growing based on propagating local constraints of reflectance and range also permits curved surfaces to be similarly grown. Attempts to differentiate between a 200 square centimeter planar surface and a 10 centimeter diameter sphere using a postprocessing step which measured the RMS error between the fitted plane and the data did not prove successful since the magnitude of range data variation was large relative to the curvature observed.

## THE TECHNICAL ARTS WHITE SCANNER

### White Scanner Operational Characteristics

The Technical Arts White Scanner 100 produces a dense three-dimensional range image by triangulating rays from the optical center of a camera through the image plane as they intersect a plane of laser light falling on objects in a scene. The three-dimensional coordinates thus produced are 16 bit values measured in increments of 0.001" (0.025 mm). Because the system relies on triangulating image rays through a plane of laser light, the system produces images that necessarily have missing data. This is because certain regions of the scene upon which the laser plane may be projected cannot be seen by the camera and certain regions visible to the camera may not be reached by the light plane. These missing data values are, however, flagged for the system user as invalid coordinates using a status byte.

### Observed Characteristics of the White Scanner

Because the images analyzed were contained in an archive of scanned images, the extensive testing on various types of surfaces with differing albedos was not possible. However, for those images that were examined, the quality of three-dimensional data was exceptionally good. Specifically, local standard deviations rarely exceeded 1 mm. Such good resolution is, of course, to be expected since the system measures in thousandths of an

inch increments. This high quality of range measurements made it possible to find most (but not all) planes in scenes even without any data smoothing, and all planes could be found by using relatively small (3X3) smoothing filters.

The most significant observable difference between the data provided by the White Scanner and the Perceptron was noted when curved objects were scanned. Curvature was fully evident in the White Scanner images containing spheres, and their local surface normals clearly indicated vectors that converged toward the center of each sphere. This permitted the reliable use of surface normals to clearly distinguish between curved and planar surfaces.

### **Teleos Prism 3 Operational Characteristics**

The Prism 3 Stereo Vision System uses two cameras to determine the distance to corresponding regions of interest in the left and right cameras. These regions of interest are processed in each image by applying a Laplacian-of-Gaussian operator and shifting (horizontally and vertically) the result of this operation in one image such that correspondence is established. There are two methods by which this window may be shifted. In the first case, the vergence angle between the two cameras may be adjusted, thus altering the fields of view for the cameras. The second method is to compute a stereo disparity by shifting the window (in software) until a match is obtained. For producing a dense range image, the second method is preferred because no mechanical repositioning of the camera system is required.

## **THE TELEOS PRISM 3 STEREO VISION SYSTEM**

### **Observed Characteristics of the Prism 3**

The Prism 3 produces the best quality of measurements when highly visually textured objects are being observed. For test purposes, two such objects, a cylinder and a planar surface, were covered with a white material upon which there was a black speckled pattern. The cylinder and planar surface were observed under varying conditions relative to each other with the following results.

When the cylinder was placed in contact with the plane at a distance of about 2 meters, the curvature of the cylinder was clearly evident, although it blended into the background plane in a manner resembling a Gaussian curve. In this configuration, the camera was both focused and verged on the cylinder.

With the cylinder remaining in the same position, the plane was moved further away from the Prism 3. As it moved further away, significant anomalies in the data developed, especially near the (jump) boundaries of the cylinder. With the background plane at a very large distance (effectively removed from the scene), the cylinder lost a significant amount of its observed curvature and the spikes near the edge of the cylinder became very large.

Although significant deterioration in data quality was noticed as the separation between the plane and cylinder increased, it was found to be possible to use internal correlation measures computed by the Prism 3 as the basis for confidences that could distinguish between good and bad data. The confidence measures were based on correlation peak strength, distinctiveness, and uniqueness. Examining images in the context of these measures revealed that although spurious measurements existed at the object boundaries, their confidences were low relative to measurements not near boundaries.

In order to obtain an estimate of RMS errors for the raw data itself, range values for the solitary plane were used. It was found that at distances of approximately 2 meters the standard deviation for range values was less than 0.8 cm.

It should be noted in all the tests for the Prism 3 that accurate focus and vergence are critical in order to obtain the best results. Hence, for the tests involving two objects (as in the cases involving the cylinder and plane), unless the lenses used have a large depth-of-focus, the range values for one object will necessarily deteriorate. Similarly, if the system is verged on the forward object, then vergence on the rear object becomes worse as it moves further away.